# EAST Search History

| Ref # | Hits | Search Query | DBs | Default Operator | Plurals | Time Stamp |
|---|---|---|---|---|---|---|
| L1 | 31 | (number with files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:18 |
| L2 | 3080 | ((backup or (back near up) or archiv$3) near database$1) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:42 |
| L3 | 0 | 1 and 2 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:42 |
| L4 | 0 | (recover$3 with (achiv$4 or backup)) and (number with files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:43 |
| L5 | 0 | (recover$3 with (achiv$4 or backup)) and (pluralit$3 with files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:43 |
| L6 | 3 | (asynchronous with files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:51 |
| L7 | 1244 | (asynchronous with (transfer$3 or transmit$4) with (simultaneously or parallel)) and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:52 |
| L8 | 3 | 7 and 2 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:52 |

# EAST Search History

| | | | | | | |
|---|---|---|---|---|---|---|
| L9 | 0 | (number with files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and backup and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:18 |
| L10 | 0 | (number with files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and achiv$3 and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:19 |
| L11 | 26 | ( files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and backup and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:19 |
| L12 | 18 | ( files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and backup and @ad<"20001206" and recover$3 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:21 |
| L13 | 3 | ( files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and backup and @ad<"20001206" and recover$3 and asynchronous | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:20 |
| L14 | 10 | ( files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and backup and @ad<"20001206" and recover$3 and predetermin$3 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:21 |
| L15 | 10 | 11 and (707/200-204).ccls. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:30 |
| L16 | 3 | 11 and (707/200-204).ccls. and asynchronous | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:31 |

# EAST Search History

| | | | | | | |
|---|---|---|---|---|---|---|
| S1 | 408 | ((backup or (back near up) or archiv$3) near database$1).ab. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:42 |
| S2 | 135 | ((backup or (back near up) or archiv$3) near database$1).ti. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:47 |
| S3 | 85 | S1 and S2 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:45 |
| S4 | 280 | ((backup or (back near up) or archiv$3) near database$1).clm. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:45 |
| S5 | 9 | S3 and S4 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:45 |
| S6 | 1 | S5 and (asynchronous or asynchroniz$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 14:54 |
| S7 | 1 | "5745753".PN. | USPAT; USOCR | OR | OFF | 2006/08/14 17:47 |
| S8 | 46 | (((backup or (back near up) or archiv$3) near database$1) and log$2).ab. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:48 |
| S9 | 0 | S8 and (asynchronous or asynchroniz$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:48 |

# EAST Search History

| | | | | | | |
|---|---|---|---|---|---|---|
| S10 | 6 | (((backup or (back near up) or archiv$3) near database$1) and log$2).ti. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:48 |
| S11 | 0 | S10 and (asynchronous or asynchroniz$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:48 |
| S12 | 68 | (((backup or (back near up) or archiv$3) near database$1) and log$2).clm. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/14 17:48 |
| S13 | 8 | S12 and (asynchronous or asynchroniz$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:07 |
| S14 | 35 | (background near synchronous) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:09 |
| S15 | 408 | ((backup or (back near up) or archiv$3) near database$1).ab. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:08 |
| S16 | 0 | S14 and S15 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:08 |
| S17 | 2823 | ((backup or (back near up) or archiv$3) near database$1) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:09 |

# EAST Search History

| | | | | | | |
|---|---|---|---|---|---|---|
| S18 | 0 | S14 and S17 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:09 |
| S19 | 727 | (background with synchronous) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:09 |
| S20 | 2 | S19 and S17 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:14 |
| S21 | 4330 | (automatic$4 near recover$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:14 |
| S22 | 82 | S21 and (707/200-204).ccls. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 12:28 |
| S23 | 2 | S22 and S15 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2006/08/15 10:15 |
| S24 | 16 | (archiv$3 or backup) with log with (asynchronous or asynchroniz$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:03 |
| S25 | 5 | S24 and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 15:50 |

# EAST Search History

| | | | | | | |
|---|---|---|---|---|---|---|
| S26 | 10 | (archiv$3 or backup) with log with streams | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:04 |
| S27 | 3 | S26 and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:03 |
| S28 | 26 | (archiv$3 or backup) with log with config$5 with file | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:09 |
| S29 | 3 | S28 and (asynchronous or asynchroniz$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:09 |
| S30 | 0 | S29 and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:10 |
| S31 | 5200 | determin$3 with (asynchronous or asynchroniz$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:04 |
| S32 | 2287 | (archiv$3 or backup) with log | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:05 |
| S33 | 529 | (log near transactions) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:05 |

# EAST Search History

| | | | | | | |
|---|---|---|---|---|---|---|
| S34 | 0 | S31 and S32 and S33 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:05 |
| S35 | 4 | S31 and S33 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:05 |
| S36 | 376 | log with (asynchronous or asynchroniz$3) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:09 |
| S37 | 66 | S32 and S36 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:09 |
| S38 | 572 | log with config$5 with file | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:10 |
| S39 | 66 | S37 and S37 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:10 |
| S40 | 26 | S39 and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 16:17 |
| S41 | 18 | S39 and @ad<"20001206" and (simultaneously or parallel) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/16 17:14 |

# EAST Search History

| | | | | | | |
|---|---|---|---|---|---|---|
| S42 | 7 | (recover$3 with (achiv$4 or backup)) and (files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:43 |
| S43 | 5 | (predetermin$4 with files with (transfer$3 or transmit$4) with (simultaneously or parallel)) and @ad<"20001206" | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | OFF | 2007/01/17 11:42 |

**Web** Results **1** - **10** of about **18** for **asynchronous backup recovery simultaneously "asynchronous stream**

[PDF] 26 Prospects for Building Highly Assured Web Services
File Format: PDF/Adobe Acrobat
structured as **asynchronous streams** (often with message queuing ... **simultaneously**. Similarly, we pointed to ways to make the server itself highly available. ...
www.springerlink.com/index/x6g2370x326765q5.pdf - Similar pages

Patentee Index
Nitrogen **recovery** system and method using heated air as stripping gas 07153427 ... in a **simultaneous** multi-threaded (SMT) processor 07155600 Cl. 712-229. ...
uspto.gov/web/patents/patog/week52/OG/patentee/alphaB.htm - 181k -
Cached - Similar pages

[PDF] SunOS Reference Manual
File Format: PDF/Adobe Acrobat
qwriter – **asynchronous STREAMS** perimeter upgrade. SYNOPSIS ... type of lock will allow many threads to have **simultaneous** read-only access to an object. ...
192.18.109.11/801-6680-9F/801-6680-9F.pdf - Similar pages

[PDF] LAWS - LGOL-Net Manual
File Format: PDF/Adobe Acrobat - View as HTML
support for defining **backup** routes for messages to provide further ... separate sources or recombining separate **asynchronous streams** of a workflowed ...
www.lgolnet.org/downloads/Documents/Current/LGOL-Net-Manual_v2_1_0.pdf - Similar pages

Current Internet-Drafts This summary sheet provides a short ...
Such devices have many voice connections **simultaneously** between them. ... "IEEE 1394 **Asynchronous Streams**", Peter Johansson, 11/17/1998, ...
quimby.gnus.org/internet-drafts/1id-abstracts.txt - Similar pages

[PDF] The Design Philosophy of Distributed Programming Systems: the ...
File Format: PDF/Adobe Acrobat
**Recovery**: In Mozart, time-outs, in the sense of when the suspended operation on ... timization, and **asynchronous streams**. For clarity, we do not define the ...
www.sics.se/~perbrand/mozart.pdf - Similar pages

[PDF] Real Time Intrusion Detection
File Format: PDF/Adobe Acrobat
multiple subnets **simultaneously**. It also cannot detect any host-based ... data traffic is ultimately passed through the **asynchronous STREAMS** stack even ...
www.ensta.fr/~hammami/DFR/Voies/MP-101-__ALL.pdf - Similar pages

@TECHREPORT{Ball9211:Core, AUTHOR="Anthony Ballardie and Paul ...
Such devices have many voice connections **simultaneously** between them. ...
TITLE="{IEEE} 1394 **Asynchronous Streams**", TYPE="Internet Draft", ...
www.cs.columbia.edu/~hgs/bib/i-d.bib - Similar pages

[PDF] Basic LAN and WAN Access Configuration Guide
File Format: PDF/Adobe Acrobat
**simultaneously**, each device must determine whether the physical medium is in ... To

synchronize the **asynchronous streams**,. the multiplexers on the line use ...
www.juniper.net/.../software/jseries/junos75/jseries75-config-guide-basic/jseries75-config-guide-basic.pdf - Similar pages

[PDF] Realtime Signal Processing
File Format: PDF/Adobe Acrobat
Modelling **Asynchronous Streams** in Haskell [115] develops Haskell code for ... evaluate these redexes **simultaneously**. Peyton Jones [104] describes the issues ...
ptolemy.eecs.berkeley.edu/~johnr/papers/pdf/thesis.pdf - Similar pages

Result Page: 1 2 **Next**

---

**Web** Results **1** - **10** of about **307** for **asynchronous streams backup recovery parallel simultaneously "arch**

## Scholarly articles for **asynchronous streams backup recovery parallel simultaneously "archive logs"**

ObjectStore Technical Overview - Design - Cited by 3

### Features and Architecture

Multiple datafiles or tablespaces can be backed up **simultaneously** to multiple devices in **parallel**. This fast, parallelized **backup** reduces the time required **...**
www.lsbu.ac.uk/oracle/oracle7/server/doc/EBADM/chap1.htm - 41k -
Cached - Similar pages

### Database **Recovery**

Using **parallel recovery**, several processes **simultaneously** apply changes from redo log files. **...** information about backups of datafiles and **archive logs ...**
www.csee.umbc.edu/help/oracle8/server.815/a67781/c28recov.htm - 86k -
Cached - Similar pages

### [PDF] OTN Case Study - Oracle Data Guard

File Format: PDF/Adobe Acrobat - View as HTML
Data Guard. LGWR **Asynchronous**. redo shipping. Disaster **Recovery** Site **...** to all database files, online logs, **archive logs** and the control file in order to **...**
www.oracle.com/technology/deploy/availability/pdf/OracleGlobalITProfile.pdf - Similar pages

### [PDF] D B & R S B P

File Format: PDF/Adobe Acrobat - View as HTML
If you cannot afford to loose data, your **backup**. plan must include the ability to **backup archive logs**. Archived redo logs are crucial for **recovery** when no **...**
www.oracle.com/technology/deploy/availability/pdf/BR_OOW01_213WP.pdf - Similar pages
[ More results from www.oracle.com ]

### [PDF] DB2 UDB **Backup** and **Recovery**

File Format: PDF/Adobe Acrobat - View as HTML
**Backup** and **recovery** utilities. DB2 UDB provides a granular and **parallel backup** and restore utility. Some of. the options available to **backup** include: **...**
www.redbooks.ibm.com/redbooks/pdfs/sg246557.pdf - Similar pages

### [PDF] Technical Report Template

File Format: PDF/Adobe Acrobat - View as HTML
supported, allowing **backup** and **recovery** to any capable system. **Backup** images are written using a. derivative of the BSD dump **stream** format, allowing full **...**
www.netapp.com/library/tr/3369.pdf - Similar pages

### [PDF] Sun Cluster 3.0 Software Cluster File System (CFS ...

File Format: PDF/Adobe Acrobat - View as HTML
simplifying the storage of these **archive logs** compared to Sun Cluster 2.2, but there **...** However, best practices for **backup** and restore are **...**
www.sun.com/software/whitepapers/wp-globalfileservices/wp-globalfileservices.pdf - Similar pages

[PDF] **Backup** and Restore Practices for Sun Enterprise™ Servers
File Format: PDF/Adobe Acrobat - View as HTML
database snapshot and feeds **parallel** data **streams** to the **backup** tool for ... **archive logs** to reduce time of **recovery**. This approach means that they are ...
docs.cirkva.net/Sun/Blue_printy/books/brbp.pdf - Similar pages

EDW Strategy
A solid **backup** and **recovery** strategy should not be limited to database file ... This approach uses an alternative form of **asynchronous** replication which ...
guweb.georgetown.edu/uis/ia/dw/edwStrategy.htm - 345k - Cached - Similar pages

[PDF] Oracle Database **Backup** and **Recovery** Advanced User's Guide
File Format: PDF/Adobe Acrobat
unique name with that **stream**. When RMAN needs to restore the **backup**, ... or **archive logs** are in **recovery** area, then stop managed **recovery** mode and ...
www.stanford.edu/dept/itss/docs/oracle/10g/server.101/b10734.pdf - Similar pages

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

---

Terms used
**asynchronous streams backup recovery parallel simultaneously archive logs**

Found **11,805** of **196,780**

Sort results by: relevance
Display results: expanded form

Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown

Relevance scale

1 Programming languages for distributed computing systems
Henri E. Bal, Jennifer G. Steiner, Andrew S. Tanenbaum
September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3
**Publisher:** ACM Press
Full text available: pdf(6.50 MB) Additional Information: full citation, abstract, references, citings, index terms, review

When distributed systems first appeared, they were programmed in traditional sequential languages, usually with the addition of a few library procedures for sending and receiving messages. As distributed applications became more commonplace and more sophisticated, this ad hoc approach became less satisfactory. Researchers all over the world began designing new programming languages specifically for implementing distributed applications. These languages and their history, their underlying pr ...

2 Experience Using Multiprocessor Systems—A Status Report
Anita K. Jones, Peter Schwarz
June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2
**Publisher:** ACM Press
Full text available: pdf(4.48 MB) Additional Information: full citation, references, citings, index terms

3 The process group approach to reliable distributed computing
Kenneth P. Birman
December 1993 **Communications of the ACM**, Volume 36 Issue 12
**Publisher:** ACM Press
Full text available: pdf(6.00 MB) Additional Information: full citation, references, citings, index terms

**Keywords:** fault-tolerant process groups, message ordering, multicast communication

4 Query evaluation techniques for large databases
Goetz Graefe
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

**Publisher:** ACM Press

Full text available: pdf(9.37 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords**: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

5 A time-sensitive object model for real-time systems

H. Rebecca Callison

July 1995 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 4 Issue 3

**Publisher:** ACM Press

Full text available: pdf(2.16 MB) Additional Information: full citation, abstract, references, citings, index terms

Process-oriented models for real-time systems focus on the timing constraints of processes, a focus that can adversely affect resulting designs. Data dependencies between processes create scheduling interactions that limit the times at which processes may execute. Processes are then designed to fit available windows in the overall system schedule. "Fitting in" frequently involves fragmenting processes to fit scheduling windows and/or designing program and data s ...

**Keywords**: concurrency, fault tolerance, object models, programming techniques, real-time processing models, timing constraints

6 Disaster recovery techniques for database systems

Manhoi Choy, Hong Va Leong, Man Hon Wong

November 2000 **Communications of the ACM**

**Publisher:** ACM Press

Full text available: pdf(412.04 KB) Additional Information: full citation, references, index terms

7 Distributed systems - programming and management: On remote procedure call

Patrícia Gomes Soares

November 1992 **Proceedings of the 1992 conference of the Centre for Advanced Studies on Collaborative research - Volume 2 CASCON '92**

**Publisher:** IBM Press

Full text available: pdf(4.52 MB) Additional Information: full citation, abstract, references, citings

The Remote Procedure Call (RPC) paradigm is reviewed. The concept is described, along with the backbone structure of the mechanisms that support it. An overview of works in supporting these mechanisms is discussed. Extensions to the paradigm that have been proposed to enlarge its suitability, are studied. The main contributions of this paper are a standard view and classification of RPC mechanisms according to different perspectives, and a snapshot of the paradigm in use today and of goals for t ...

8 The TickerTAIP parallel RAID architecture

Pei Cao, Swee Boon Lin, Shivakumar Venkataraman, John Wilkes
August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3
**Publisher:** ACM Press

Full text available: pdf(2.04 MB) Additional Information: full citation, abstract, references, citings, index terms

Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the maximum number of disks to which the array can scale. We describe TickerTAIP, a parallel architecture for disk arrays that distributes the controller functions across several loosely coupled processors. The result is better scalability, fault tolerance, and flexibility.This article present ...

**Keywords**: RAID disk array, decentralized parity calculation, disk scheduling, distributed controller, fault tolerance, parallel controller, performance simulation

9 Cluster communication protocols for parallel-programming systems

Kees Verstoep, Raoul A. F. Bhoedjang, Tim Rühl, Henri E. Bal, Rutger F. H. Hofman
August 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 3
**Publisher:** ACM Press

Full text available: pdf(1.29 MB) Additional Information: full citation, abstract, references, index terms, review

Clusters of workstations are a popular platform for high-performance computing. For many parallel applications, efficient use of a fast interconnection network is essential for good performance. Several modern System Area Networks include programmable network interfaces that can be tailored to perform protocol tasks that otherwise would need to be done by the host processors. Finding the right trade-off between protocol processing at the host and the network interface is difficult in general. In ...

**Keywords**: Clusters, parallel-programming systems, system area networks

10 Distributed logging for transaction processing

Dean S. Daniels, Alfred Z. Spector, Dean S. Thompson
December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87**, Volume 16 Issue 3
**Publisher:** ACM Press

Full text available: pdf(1.51 MB) Additional Information: full citation, abstract, references, citings, index terms

Increased interest in using workstations and small processors for distributed transaction processing raises the question of how to implement the logs needed for transaction recovery. Although logs can be implemented with data written to duplexed disks on each processing node, this paper argues there are advantages if log data is written to multiple log server nodes. A simple analysis of expected logging loads leads to the conclusion that a high performance, microprocessor b ...

11 A taxonomy of Data Grids for distributed data sharing, management, and processing

Srikumar Venugopal, Rajkumar Buyya, Kotagiri Ramamohanarao
June 2006 **ACM Computing Surveys (CSUR)**, Volume 38 Issue 1
**Publisher:** ACM Press

Full text available: pdf(1.70 MB) Additional Information: full citation, abstract, references, index terms

Data Grids have been adopted as the next generation platform by many scientific communities that need to share, access, transport, process, and manage large data collections distributed worldwide. They combine high-end computing technologies with high-performance networking and wide-area storage management techniques. In this article, we discuss the key concepts behind Data Grids and compare them with other data sharing and distribution paradigms such as content delivery networks, peer-to-peer n ...

**Keywords**: Grid computing, data-intensive applications, replica management, virtual organizations

12 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

**Publisher:** ACM Press

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

13 Providing high availability using lazy replication

Rivka Ladin, Barbara Liskov, Liuba Shrira, Sanjay Ghemawat

November 1992 **ACM Transactions on Computer Systems (TOCS)**, Volume 10 Issue 4

**Publisher:** ACM Press

Full text available: pdf(2.46 MB) Additional Information: full citation, abstract, references, citings, index terms, review

To provide high availability for services such as mail or bulletin boards, data must be replicated. One way to guarantee consistency of replicated data is to force service operations to occur in the same order at all sites, but this approach is expensive. For some applications a weaker causal operation order can preserve consistency while providing better performance. This paper describes a new way of implementing causal operations. Our technique also supports two other kinds of operations: ...

**Keywords**: client/server architecture, fault tolerance, group communication, high availability, operation ordering, replication, scalability, semantics of application

14 Reliability mechanisms for SDD-1: a system for distributed databases

Micael Hammer, David Shipman

December 1980 **ACM Transactions on Database Systems (TODS)**, Volume 5 Issue 4

**Publisher:** ACM Press

Full text available: pdf(3.06 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents the reliability mechanisms of SDD-1, a prototype distributed database system being developed by the Computer Corporation of America. Reliability algorithms in SDD-1 center around the concept of the Reliable Network (RelNet). The RelNet is a communications medium incorporating facilities for site status monitoring, event timestamping, multiply buffered message delivery, and the atomic control of distributed transactions. This paper is one of a series of compani ...

**Keywords**: atomicity, distributed databases, recovery, reliability

15 MPICH-V2: a Fault Tolerant MPI for Volatile Nodes based on Pessimistic Sender Based Message Logging

Bouteiller Bouteiller, Franck Cappello, Thomas Herault, Krawezik Krawezik, Pierre Lemarinier, Magniette Magniette
November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing SC '03**
**Publisher:** IEEE Computer Society
Full text available: pdf(527.20 KB) Additional Information: full citation, abstract, citings

Execution of MPI applications on clusters and Grid deployments suffering from node and network failures motivates the use of fault tolerant MPI implementations. We present MPICH-V2 (the second protocol of MPICH-V project), an automatic fault tolerant MPI implementation using an innovative protocol that removes the most limiting factor of the pessimistic message logging approach: reliable logging of in transit messages. MPICH-V2 relies on uncoordinated checkpointing, sender based message logging ...

16 A Teradata content-based multimedia object manager for massively parallel architectures

W. O'Connell, I. T. Ieong, D. Schrader, C. Watson, G. Au, A. Biliris, S. Choo, P. Colin, G. Linderman, E. Panagos, J. Wang, T. Walter
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2
**Publisher:** ACM Press
Full text available: pdf(1.18 MB) Additional Information: full citation, abstract, citings, index terms

The Teradata Multimedia Object Manager is a general-purpose content analysis multimedia server designed for symmetric multiprocessing and massively parallel processing environments. The Multimedia Object Manager defines and manipulates user-defined functions (UDFs), which are invoked in parallel to analyze or manipulate the contents of multimedia objects. Several computationally intensive applications of this technology, which use large persistent datasets, include fingerprint matching, signatur ...

**Keywords**: Teradata, content-based analysis, parallel multimedia database, user-defined functions

17 Conference abstracts

January 1977 **Proceedings of the 5th annual ACM computer science conference CSC '77**
**Publisher:** ACM Press
Full text available: pdf(3.14 MB) Additional Information: full citation, abstract, index terms

One problem in computer program testing arises when errors are found and corrected after a portion of the tests have run properly. How can it be shown that a fix to one area of the code does not adversely affect the execution of another area? What is needed is a quantitative method for assuring that new program modifications do not introduce new errors into the code. This model considers the retest philosophy that every program instruction that could possibly be reached and tested from the ...

18 A coherent distributed file cache with directory write-behind

Timothy Mann, Andrew Birrell, Andy Hisgen, Charles Jerian, Garret Swart
May 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 2
**Publisher:** ACM Press
Full text available: pdf(3.21 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Extensive caching is a key feature of the Echo distributed file system. Echo client machines maintain coherent caches of file and directory data and properties, with write-behind (delayed write-back) of all cached information. Echo specifies ordering constraints

on this write-behind, enabling applications to store and maintain consistent data structures in the file system even when crashes or network faults prevent some writes from being completed. In this paper we describe ...

**Keywords**: coherence, file caching, write-behind

**19** DVM: an object-oriented framework for building large distributed Ada systems

Christopher J. Thompson, Vincent Celier

November 1995 **Proceedings of the conference on TRI-Ada '95: Ada's role in global markets: solutions for a changing complex world TRI-Ada '95**

**Publisher:** ACM Press

Full text available: pdf(1.50 MB) Additional Information: full citation, references

**20** Rx: treating bugs as allergies---a safe method to survive software failures

Feng Qin, Joseph Tucek, Jagadeesan Sundaresan, Yuanyuan Zhou

October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5

**Publisher:** ACM Press

Full text available: pdf(245.29 KB) Additional Information: full citation, abstract, references, citings, index terms

Many applications demand availability. Unfortunately, software failures greatly reduce system availability. Prior work on surviving software failures suffers from one or more of the following limitations: Required application restructuring, inability to address deterministic software bugs, unsafe speculation on program execution, and long recovery time.This paper proposes an innovative *safe* technique, called Rx, which can quickly recover programs from many types of software bugs, both det ...

**Keywords**: availability, bug, reliability, software failure

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Results for "((asynchronous streams backup recovery parallel simultaneously 'archive logs')<in>metadata)" e-mail
Your search matched 0 documents.
A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

View Session History

New Search

Modify Search

((asynchronous streams backup recovery parallel simultaneously 'archive logs')<in>m Search

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

| | |
|---|---|
| IEEE JNL | IEEE Journal or Magazine |
| IEE JNL | IEE Journal or Magazine |
| IEEE CNF | IEEE Conference Proceeding |
| IEE CNF | IEE Conference Proceeding |
| IEEE STD | IEEE Standard |

**No results were found.**

Please edit your search criteria and try again. Refer to the Help pages if you need assistan search.